

AD-A218 886

**SENTENCE COMPREHENSION: A
PARALLEL DISTRIBUTED PROCESSING
APPROACH**

Technical Report AIP - 70

**James L. McClelland Mark St. John
& Roman Taraban**

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

**The Artificial Intelligence
and Psychology Project**

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

Approved for public release; distribution unlimited.

90 03 12 030

DTIC
S **ELECTE** **D**
MAR 13 1990
B

2

SENTENCE COMPREHENSION: A PARALLEL DISTRIBUTED PROCESSING APPROACH

Technical Report AIP - 70

**James L. McClelland Mark St. John
& Roman Taraban**

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

14 July 1989

**DTIC
ELECTE
MAR 13 1990
S B D**

The work reported here was supported by NSF Grants BNS 86-09729 and BNS 88-12048, ONR Contracts N00014-86-0146 and N00014-86-K-0349, and an NIMH Career Development Award (MH00385) to the first author.

This research was supported by the Computer Sciences Division, Office of Naval Research, under contract number N00014-86-K-0678. Reproduction in whole or part is permitted for any purpose of the United States Government. Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE					
1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; Distribution unlimited		
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) AIP - 70			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University		6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research		
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213			7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization		8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678		
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS p4000ub201/7-4-86		
			PROGRAM ELEMENT NO. N/A	PROJECT NO. N/A	TASK NO. N/A
			WORK UNIT ACCESSION NO. N/A		
11. TITLE (Include Security Classification) Sentence comprehension: A parallel distributed processing approach					
12. PERSONAL AUTHOR(S) McClelland, James L., St. John, Mark and Taraban, Roman					
13a. TYPE OF REPORT Technical		13b. TIME COVERED FROM 86Sept15 TO 91Sept14		14. DATE OF REPORT (Year, Month, Day) 1989 July 14	
				15. PAGE COUNT 53	
16. SUPPLEMENTARY NOTATION In press, Language and Cognitive Processes					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP			
			Neural networks; connectionist models; language comprehension; language acquisition.		
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
SEE REVERSE SIDE					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION		
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Mayrowitz			22b. TELEPHONE (Include Area Code) (202) 696-4302		22c. OFFICE SYMBOL N00014

Abstract

We review basic aspects of conventional approaches to sentence comprehension and point out some of the difficulties faced by models that take these approaches. We then describe an alternative approach, based on the principles of parallel distributed processing, and show how it offers different answers to basic questions about the nature of the language processing mechanism. We describe an illustrative simulation model that captures the key characteristics of the approach, and illustrates how it can cope with the difficulties faced by conventional models. We describe alternative ways of conceptualizing basic aspects of language processing within the framework of this approach, consider how it can address several arguments that might be brought to bear against it, and suggest avenues for future development. *Keyword.*



Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

What is constructed mentally when we comprehend a sentence? How does this constructive process occur? What role do words play in the construction process? How is the ability to construct such a representation acquired? These are some of the the central questions that face any attempt to build a model of language processing.

In this paper, we present a view that differs from some existing notions about the general form of the answers to these questions. We briefly outline what we take to be a generic version of existing notions. Then, we point out some difficulties with these notions. After this, we present a sketch of an alternative. We illustrate the alternative with a preliminary model, and consider how it gives different answers to some of the questions raised above. We examine some of the arguments, both theoretical and empirical, that have been taken as counting against this sort of alternative. Finally, we describe future directions for the further development of this approach.

1. Conventional approaches to sentence comprehension.

The comprehension of sentences has of course been studied extensively, and there are many disparate views about the nature of this process. We do not mean to assert that all previous researchers have adhered to the views we describe in this section. However, quite a bit of work has been done which we believe either tacitly or explicitly adopts the views we describe here. We tend to cite the paper by Fodor and Pylyshyn (1988), because it articulates these views extremely clearly. Where relevant, we will site works that apply these ideas and general texts where they are used or assumed.

1.1. What is constructed when we comprehend a sentence? It is typical to assume that what is constructed is an interconnected set of propositions (e.g., Clark and Clark, 1977), or *propositional representation*. The exact nature of these propositions varies from implementation to implementation, but in general they are taken to be symbolic expressions which have a combinatorial syntax and semantics (Fodor & Pylyshyn, 1988). According to Fodor and Pylyshyn, combinatorial representations are those which exhibit the following

properties:

- They may be atomic or molecular expressions.
- If they are molecular, they have constituents which may be either atomic or molecular.
- The semantic content of a molecular expression is a function of the semantic content of each of the parts of the expression and of the organization of the constituents.

1.2. What role do words play in the comprehension process? Implicit in many theories of comprehension is the notion that words have meanings, and that these meanings are the constituents of the meanings of the propositions that are constructed from sentences that contain these words. This view appears to underlie Fodor and Pylyshyn's (1988) principle of *compositionality*: According to this principle, "a word makes approximately the same semantic contribution to the meaning of every sentence in which it occurs." Let us use their example:

1. John loves the girl.
2. The girl loves John.

Fodor and Pylyshyn use these sentences to illustrate what they mean by compositionality. They ask us to consider the meaning of the word "loves" that appears in both of these sentences. They state that the relationship that John is said to bear to the girl in the first sentence is the same relationship that the girl is said to bear to John in the second sentence. This common relationship can be taken to be the meaning of the word "loves", and it occurs in the representation of the meaning of both of these sentences.

1.3. How does the process of constructing a representation of the propositions underlying a sentence occur? Often, this process is taken to be one of building a structural description using a system of *structure sensitive* rules. Following Fodor and Pylyshyn, we take *structure sensitive* to mean that the operations that apply to representations are sensitive to their form and not their content (Fodor and Pylyshyn, 1988). Examples of models that attempt to characterize the construction of structural descriptions of sentences using such rules are Marcus' deterministic parser (Marcus, 1980; but see below), and

Frazier's account of initial parsing strategies (Frazier, 1986).

1.4. How is the ability to construct a representation acquired? To the extent that we assume that the process of constructing representations of sentences proceeds by the use of structure-sensitive rules to structure the constituent expressions corresponding to words, it seems natural to assume that acquisition amounts to a process of determining what the rules are and what the constituent expressions are that words are used to designate. Researchers interested in acquisition of comprehension skill do not of course assume that the rules that are actually used in comprehension are the same rules that characterize the abstract linguistic competence of the speaker-hearer, but this does not mean that such rules are not rules nonetheless.

1.5. Summary. In brief, the comprehension of sentences is generally taken to be the process whereby a listener uses a set of structure sensitive rules to construct a propositional representation that constitutes the "meaning" of the sentence. The constituents of this representation include the meanings of the words in the sentence. Following Fodor and Pylyshyn's terminology, we call this view the *classical view*. These authors intend it to be taken as applying more broadly than to just the interpretation of sentences, but they make clear that language is a "paradigm of systematic cognition"; we will not have anything to say about its broader applicability; instead we will focus on the reasons why we feel that this view is not applicable specifically to language comprehension.

2. Problems for the classical view of sentence comprehension.

2.1. Conceptual guidance and rule conflicts. A central problem for the conventional view is the fact that sentence interpretations cannot in general be recovered correctly from structure sensitive rules alone. Even those who try to go the farthest using structure sensitive rules (Marcus, 1980; Frazier, 1986) are accurately aware of this problem. The problem is not just a curiosity; it comes up most every time a prepositional phrase is encountered. Consider:

3. The spy saw the policeman with binoculars.
4. The spy saw the policeman with a revolver.

In 3, most readers interpret the binoculars as the instrument used by the spy in seeing the policeman. In 4, most readers interpret the revolver as a possession of the policeman. This simple example illustrates clearly that it is necessary at a minimum to consider whether the object of the prepositional phrase is a plausible candidate for use as an instrument of the verb. In general, as the next example makes clear, it is also necessary to consider whether in fact the agent of the sentence might be the kind of agent that can use the instrument:

5. The bird saw the birdwatcher with binoculars.

Indeed, Oden (1978) has shown that every constituent of sentences like 3-5 can potentially influence the interpretation of the role of the noun-phrase.

It is widely accepted that the *ultimate* interpretation that a sentence receives is affected by content. Many researchers accept this, but resist the idea that the initial processing of attachment ambiguities is influenced by content. Thus, for example, Frazier (1986) has proposed that initial parsing decisions are based on a purely syntactic mechanism that proposes its preferred alternative for consideration by semantic processes. Later in the paper we review empirical evidence relevant to this claim. For the moment we point out a more conceptual problem with it. The difficulty is that the decision as to which interpretation of an ambiguous sentence will win out in the end does not seem in general to be based on a simple yes-no decision about the acceptability of the supposedly syntactically preferred interpretation. Thus in 5, it is not really plausible to argue that the interpretation in which the bird is using the binoculars as instrument is strictly blocked. For example we have no difficulty accepting such an interpretation in "The bird saw its prey with binoculars", even if we find it somewhat odd for a bird to be using an instrument. Rather it appears that the alternative interpretation is simply more plausible in the case of 5. It thus appears that more than one alternative interpretation must be evaluated for plausibility, thereby robbing the parser of any special role in

providing a single alternative for consideration.

It is also important to note that it is not simply the case that decisions can either be made by syntactic rule or need to be left for semantic determination. As Marcus (1980) points out, language comprehenders have preferences for syntactic interpretation which must be seen as matters of degree, so they sometimes win and sometimes lose when placed in conflict with other considerations. Very clear examples of this arise in sentences like 6 and 7:

6. We ate some food with some friends that we like.

7. We found a painting in the attic that was covered with cobwebs.

A structure-sensitive rule would allow us to correctly parse 6, based on the idea that relative clauses should be taken to attach to the immediately preceding noun phrase rather than an earlier one, especially when, as in this case, attachment to the earlier noun-phrase would violate the so-called "no-crossover" constraint. However, it is exactly this constraint that is violated in 7, where it is the painting, rather than the attic, which native speakers take to have been covered with cobwebs. Violating this constraint may make the sentence seem a bit awkward but it does not prevent the cobwebs from attaching to the painting.

2.2. Contextual shading as well as selection of word meaning. The problem of word-meaning indeterminacy also poses a problem for conventional approaches. It is, of course, typical to assume that an individual word can have more than one meaning. The problem of sentence interpretation then is seen as one of selecting the right meaning from a set of possible meanings that are stored in a "mental lexicon". One problem with this is the potential combinatorial explosion that can result, as discussed below in 2.4. Here we focus on a different problem: The problem is that it seems rather restricting to suppose that the range of meanings that a word can have is restricted in advance to the set of known usages of the word. Let us consider some examples.

8. The hostess threw the ball for charity.

9. The slugger hit the ball over the fence.

10. The baby rolled the ball to her daddy.

The distinctions among the meanings of ball as it appears in 8 and 9 seem well enough captured by the idea that the specification of a meaning for this word involves a selection of one of two alternatives, one that means something like "fancy dance" and one that means something like "spherical object". But in 10, it seems that the specification of the ball is somewhat different from the specification that we get from 9. It is possible to assert that here again we are selecting between two alternative meanings, one, let us say, in which the spherical object is smallish, hard and white and the other in which it is larger, squishier, and probably multi-colored; but taken to its extreme, this view seems to lead to a vast explosion of lexical entries, one for each of the possible balls that we can envision being implicitly described in a sentence. Is there to be a separate lexical entry for every shade of meaning that can be comprehended, for every word in the language.

2.3. A similar problem with roles. A similar problem arises when we attempt to specify the set of structural roles that are available to be filled by word meanings in the structural description that represents a sentence. In early work on roles (Fillmore, 1968), attempts were made to enumerate the set of roles that constituents could fill. However, this effort quickly ran into the problem that there are a large number of slight distinctions among roles all of which have interpretive significance. The problem is so bad that many workers have taken the tack of assuming that for each verb there is an idiosyncratic set of roles. This is of course not terribly satisfactory either since this simply obscures the broad commonality that does exist among, for example, the constituents which we would tend to call agents if we did not look too closely.

2.4. Implied constituents. The notion that the representation of a sentence consists of an assemblage of representations of constituents of a sentence fails to provide any direct way of understanding why it is that many sentences convey implied constituents which native speakers do not need to hear mentioned. Thus in 11 and 12

11. The boy spread the jelly on the bread.
 12. The man stirred his coffee.

we can infer a knife and a spoon respectively. That such inferred constituents are expected to be parts of the representations we form in listening to sentences is indicated by the fact that we can refer to them as though they have been mentioned. Thus we can say for example

13. The boy spread the jelly on the bread.
 The knife was covered with poison.

and we can expect the reader to know that someone is in danger of being poisoned if they eat the sandwich.

Now, typically, it would be conventional to assume either that implied constituents are parts built into the representations of the lexical items (e.g., the knife is built into the representation of the verb *spread*) or that they are inferred by post-processes. However, it is by no means an easy task to decide when something should be built in; nor is it easy to decide when something *should be inferred*. We don't *always* stir coffee with a spoon, and we don't even *necessarily* spread jelly with a knife; so drawing an inference in an all-or-nothing way can lead to overcommitment. We might draw inferences and assign them strengths, but there is no end to the inferences that we might draw. Should we draw all of them? Where should the line be drawn? These problems have plagued inference based comprehension programs for years (Shank, 1981).

2.5. Combinatorial explosion or premature commitment? The multiplicity of alternative meanings of words and of possible roles, and the wide range of possible inferences which might follow from each possible combination of roles and meanings becomes an extremely serious problem when we consider the implications for processing. Famous examples like

14. Time flies like an arrow.

remind us of the potential combinatorial explosion associated with the multiplicity of possible word-meaning and structural possibilities that arise in processing virtually every sentence. Models built in the classical tradition are

forced to take one of two approaches to this problem: Either they can create a potentially exponential number of possible interpretations or they can make an early commitment to pursue only a limited range of alternatives. In the extreme form, a single track is chosen, subject to backtracking if that track turns out to fail.

The fact that native speakers are sometimes garden-pathed has often been taken as support for the view that we generally follow a single track. It seems likely that such a commitment really does occur in sentences like Bever's 15. The horse raced past the barn fell.

However, the same sense of surprise and incomprehension followed by reorganization does not occur with all ambiguities. Thus consider 16 and 17:

16. The bat flew out of the hitter's hand and hit a spectator in the stands.

17. The bat flew out of the cave and into the moonlit night.

At least for North Americans, the bat in 16 is a very different kind of object from the bat in 17. Yet no strong garden-path effect is felt in either case. Thus it would appear that the strong garden-path effect in 15 should not be taken as an indication that we always commit prematurely as we process sentences from left to right. Instead it would appear that we are able to keep a variety of options open and to use both prior and subsequent context in disambiguation.

2.6. The difficulty of acquisition. As a final note, we remind the reader of the problem of acquisition. Several serious problems face anyone who attempts to build a model of acquisition of the rules and word meanings posited by the classical view:

- The rules are often over-ridden, as we saw in 2.1.
- The possible set of rules that might be used is drastically underdetermined by the evidence available to the child.
- A given sentence may have more than one perfectly acceptable interpretation. This makes it hard to know when to reject a rule as wrong or simply not always right.
- Correct performance requires not only the knowledge of the constraints but how much weight each one should be given.
- The child faces a very serious boot-strapping problem in learning

to map sentences onto their meanings. This problem is reviewed by Gleitman and Wanner (1982).

It would prevent us from getting on with the business of this paper to explore each of these problems in detail. For now, then, we will elaborate only on the last mentioned since it will be directly addressed below.

The problem is as follows. Suppose a child hears someone say "The boy is kissing the girl." And at the same time he sees one child kissing another. Before he knows the rules of syntax, it is hard to use this sentence-event pair to know which child he should take to be the boy and which the girl. At the same time, before he knows the meanings of the words, it is hard to use this pairing to learn about the syntax. It could be, for example, that we use the word "boy" in English to refer to girls and the word "girl" to refer to boys; and that we use object-verb-subject order in describing events.

This and other problems have led many psycholinguists to the view that acquisition is impossible. Instead it has often been proposed that the rules of all languages are innate and that acquisition simply amounts to setting parameters where there are degrees of freedom. It has even been proposed (e.g., Chomsky, 1988) that it is not implausible to imagine that all concepts are innate.

2.7. Summary. We do not wish to make light of classical models. Such models do have considerable appeal, and they seem to us to capture approximately some of the general characteristics of natural languages. Indeed there are regularities in the way we structure sentences which give clues to the ideas we wish these sentences to convey; and there are regularities in the ways in which we use words. These two facts seem consistent with the idea that words have meanings that are parts of the meanings of the sentences that they occur in and that the meanings of the wholes are constructed from these parts by structure sensitive rules. Fodor and Pylyshyn (1988) are of course correct when they point to the productivity and systematicity of language, and it is no

mean accomplishment of the classical view that it captures these essential characteristics of natural language.

But it is our view that the classical approach is destined to remain strapped with all the problems listed above. It is not from any lack of appreciation of the accomplishments of classical approaches that we seek an alternative. It is only our belief that it may be possible to develop an alternative which may ultimately prove to be even more successful. The rest of this paper is an attempt to give the reader a sense of what this alternative may be like.

3. A PDP Alternative

3.1. Denied Presuppositions. The PDP alternative which we will propose denies the point of departure, implicit in classical approaches, that it is necessary to require information to be displayed in structured form in the representation itself (van Gelder, in press). Rather, we ask only that the representations provide a sufficient basis for performing the task or tasks that are required of them. Thus, representations of sentences are not required to exhibit a specifically propositional format *so long as they can be used to perform the tasks we require*. Similarly, representation of knowledge about how to form representations is not required to take the form of rules *as long as this knowledge allows us to act in lawful ways as the environment demands*, and representations of word-specific knowledge is not required to have any visible internal structure representing the meaning of the word. Indeed, the knowledge of rules and of word-specific information may well be encoded in a densely compiled form, as long as this information can be used effectively to meet the imposed demands.

3.2. Nature of the task. Our first step, then, must be to develop some conception of the nature of the imposed demands. At a general level, we think it is reasonable to think of the sentence comprehension task in the following terms. A sequence of words is presented, and the listener must form a representation which allows him to respond correctly when probed in various ways. In general, the probes can take a wide range of different forms, requiring

actions, verbal responses, etc. Among the things we would expect is that we would be able to answer various questions using this representation. For example, on hearing "The man stirred the coffee", we would expect a device that has understood this sentence to be able to give correct answers to many questions. Who did the stirring? What did he stir? What did he stir with? etc.

Given this conception of comprehension, we will need a model which can actually listen to a sentence and then respond correctly to a set of probes. Since we do not stipulate exactly what form the representations must take, we must rely on the adequacy of the performance of the model to determine if in fact its representations are adequate.

For the purposes of what follows, we will distinguish between the process of comprehension itself -- the formation of a representation from a sentence -- and the use of this representation to respond appropriately to probes. Our main interest is in the former, but for the reasons just given the latter must be considered as well or we have no measure of successful performance.

3.3. Constraint Satisfaction Processing. We think of the process of comprehension as a constraint satisfaction process (Rumelhart, Smolensky, McClelland and Hinton, 1986). In the comprehension of isolated sentences, there are two sorts of constraints: Those imposed by the sequence of words, and those imposed by knowledge about how such sequences are to be interpreted. Both types of constraints are taken to be *graded*. They are assumed to act as forces shaping the formation of a representation, and to have magnitudes which determine their degree of influence. For our purposes, the sequence of words in the sentence can be instantiated as a sequence of patterns of activation over a set of processing units. As each new word comes in, we assume that it is used to update the sentence representation, which is also taken to be a pattern of activation over a set of processing units. In fact, if we consider the process at each time step, it is useful to view it as a constraint satisfaction process in which there are two inputs: The sentence

representation from the previous time step, and the new input. These two inputs are used to produce an updated sentence representation for the next time step. The knowledge of how this updating is to be performed is stored in the connections that allow these inputs to update the sentence representation.

After each update of the sentence representation, it can be used to respond to one or more probes. Responding to these probes is also viewed as a constraint satisfaction process, where the goal is to produce externally-specified outputs in response to externally-provided probes. There are now three sources of constraint: The sentence representation, the probe, and knowledge about what outputs should be produced for particular sentence/probe combinations. Both the sentence representation and the probe can be instantiated as patterns of activation over processing units, as can the desired outputs; and the knowledge of how to produce these outputs from the corresponding inputs can be encoded in the connections among the processing units.

So far we have outlined a general framework for sentence comprehension and for using the results of comprehension to respond to probes. A sketch of the network that instantiates this framework is shown in Figure 1. In the figure, the ovals correspond to pools of units and the arrows correspond to connections. There is a pool of units for representing the successive words; a pool of units for representing the evolving sentence representation, or *Sentence Gestalt*; a pool for representing probes, and a pool for representing responses to the probes. The arrows represent connections, from each unit in the pool at the sending end of the arrow to each unit in the pool at the receiving end. The unlabelled pools of units serve to allow *combinations* of aspects of the patterns on the input side of these pools to constrain the patterns of activation on the output side.

3.4. Learning by Connection Adjustment. Three crucial questions remain. First, what determines the form of the sentence representation itself?




Figure 1: An sketch of the present conception of the sentence comprehension mechanism. The ovals represent groups of units, and the arrows represent modifiable connections.

Second, how is the form of this representation communicated to the inner part of the network? Third, how is the knowledge acquired that governs the construction of the sentence representation from the sequence of words, and the production of appropriate outputs to sentence/probe combinations? The answer to all of these questions is the same: Connection strength adjustment through error-correcting learning.

We assume that the output pattern actually generated by the network in response to each probe is compared to the correct output that is provided as part of the environment. A statistic called cross-entropy (Hinton, 1987), representing the degree to which the information content of the target is actually captured by the network's output, can be computed. Each connection weight in the network can be adjusted so as to reduce the amount of information not captured. This connection adjustment process occurs for connections in both the comprehension network and for connections in the readout network. Once this adjustment process is successful, then it has basically caused the network to discover just what the sentence representations need to look like in order to successfully constrain the generation of output in response to the probe.

4. A Model Illustrating the Approach

The model we describe here exemplifies the approach described above. It is in many ways highly simplified. It will not convince the reader that we have already succeeded in providing a complete alternative to conventional approaches. Rather, it provides an concretization of the general approach as well as an illustration of some of the reasons for its appeal which we hope will suggest that the further exploration of a new framework is worthwhile. The model is called the Sentence Gestalt or SG model. It is described briefly here; a fuller description is available in St. John and McClelland (in press).

4.1. The environment. The model consists of a network placed in an environment consisting of sentence / event-description pairs. The sentences are of but one clause, and they consist of a sequence of stripped down

constituents. Each constituent consists of a single contentive (noun, verb or adverb) together with a single preposition or the verbal auxiliary element "was". For example, the English sentence "The school girl was kissed by the boy" is reduced to three constituents, "schoolgirl", "was kissed", "by boy". The most complex sentences involved dative passives like "The teacher was given a rose by the busdriver", with additional locative, manner, and/or instrumental prepositional phrases possible depending on the verb.

The event descriptions are simple too; they consist only of a list of role-filler pairs. The roles are agent, action, object, recipient, location, manner, instrument and what might best be called "accompanist" (as in "the busdriver ate an ice-cream cone with the schoolgirl.")

While the sentences and the events they describe are both quite simple, the relationships which hold between them are not. For one thing, words used in a sentence may be ambiguous or vague, as in

- 18. The pitcher hit the ball with the bat.
- 19. The adult ate something.

In both cases, the model is asked to do its best to recover the correct event description. In the latter case, the event description involves a specific adult and a specific something eaten, which may not be uniquely predictable (in the small world of the model, the adult might be teacher or a busdriver; the something might be soup or steak). The model must do its best based on the information given.

Constituents may also be left out of sentences, as in

- 20. The busdriver stirred the coffee.

Here the network is expected to understand that the event being described involved an instrument, which in the case of stirring is always a spoon.

Role assignment is made difficult in two ways. First, both active and passive constructions are used. Though there are semantic constraints that often make correct interpretation of passives possible, this is not always the case, as

in sentences like

21. The teacher was given the rose by the busdriver.

(Note that the model does not distinguish "gave" from "given" since for most verbs the past and past participle are the same in English).

The other source of role assignment difficulty arises from the ambiguity of surface role cues. Prepositions and word order information provide some cues, but these cues are often quite ambiguous as to the roles that they signify.

Thus in

22. The busdriver ate the steak with the teacher.

23. The busdriver ate the steak with the knife.

the semantics of the role-filler must be considered in determining whether the object of the with-phrase is an instrument or an accompanist.

The actual set of sentence-event pairs that the model sees is generated as follows. First, an action is selected at random from a set of possible actions. Then, an agent is selected from a set of possible agents who might perform the action. Following this, an object, an instrument or indirect object if applicable, and other roles are filled. An illustration is given for the action "eat" in Figure 2. Note that the selection process is inherently probabilistic, and that there are complex dependencies. Given, for example, that the action is eat and the agent is busdriver, the object is probably steak ($p = .875$) but may be soup ($p = .125$); the instrument depends on the object eaten, the manner on the agent of eating.

This procedure produces an ensemble of event descriptions which are strongly constrained. These constraints can be absolute or *hard*, so that, for example, animal bats do not show up at all as the instruments of hitting; or they may be *soft*, so that, for example, steak is the preferred but not the unique object of eating for the busdriver. Note that the constraints are fairly complex, in that they depend on particular conjunctions of verbs and role fillers. Steak is the preferred food only of the busdriver, the knife is the instrument of eating when the food is steak but not soup, etc.

Structure of Events

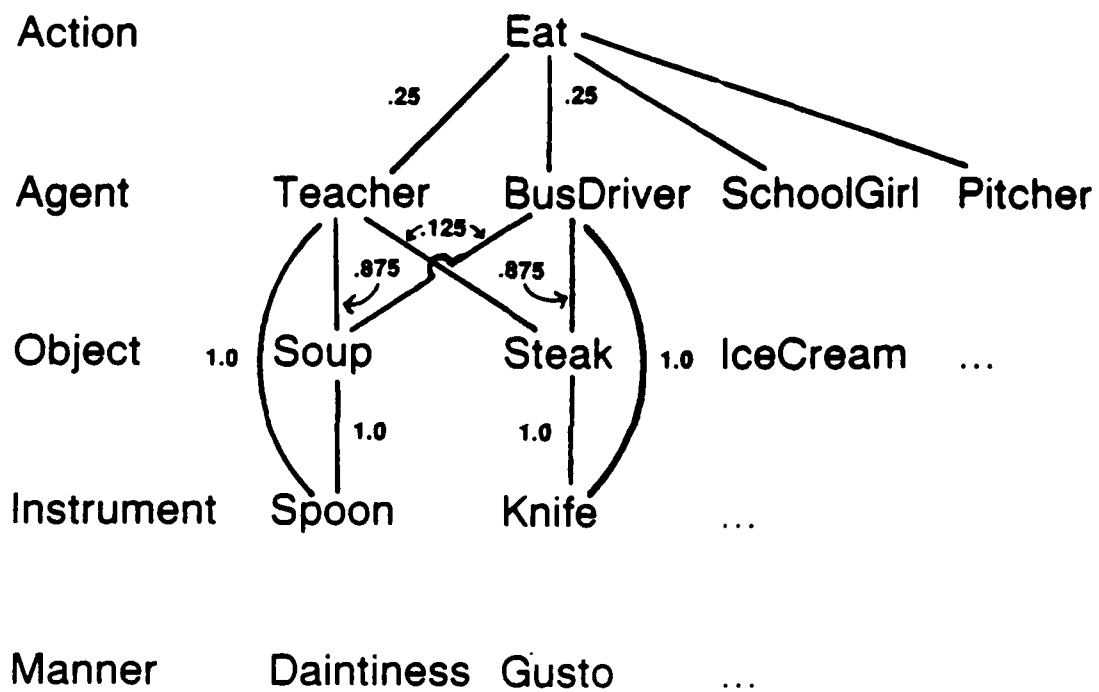


Figure 2: Structure of the event generator for the action *eat* used in training the SG model

Assignment of words to events is also probabilistic. Thus, the busdriver in the eating example might be described with the word busdriver or with the word adult; the steak might be described as steak or as food; the instrument as knife or as utensil. Thus, the actual specific participants in the events can only be inferred by using information from context. Sometimes, the sentence contains sufficient information to remove all uncertainty with respect to a particular participant (As in "the busdriver ate the food with the knife"; the food can only be steak), but other times not (as in "the busdriver ate the food"). Even here some answers may be more likely than others, though in some cases there may be at least two equally likely alternatives (in "the adult ate the food" soup and steak are equally likely).

Sometimes, whole constituents are simply left out of sentences describing events in which their referents appear. Thus, the knife can be left out of the sentence on the busdriver eating steak. The model adheres to the conventions that subject, and verb are always mentioned (however vaguely), but other constituents may go unmentioned, depending on the specific actions.

4.2. The task, and the interface to the environment. The model's task is to process the sequence of constituents that represents a particular sentence and, as each constituent comes in, to update a representation which is intended to allow it to respond to probes querying its comprehension of the event described by the sentence. To assess the model's performance, we can actually probe it after each constituent has been processed.

Each input constituent consists of a content word and possibly a preposition or "was". Each such word is represented by a single unit. Thus there is a unit for "bat" (regardless of meaning), a unit for "gave", a unit for "adult", a unit for "was", "with", "by", etc. Altogether there were units for 58 words.

A similar localist representation scheme was also used for probes and responses. Responding to a probe can be thought of as completion: filling in a member of a role-filler pair, when probed with either the role or the filler. Note

that the fillers are now concepts rather than words, and that fillers in particular events are always specific concepts, rather than superordinate categories. There were a total of 45 concepts units, covering actions, manners, and noun-concepts, including persons, places, and things.

In some simulations using this model St. John and McClelland included a few units representing superordinate concepts in addition to the units for specific concepts. In this case, a concept is represented not by a single unit, but by a set of units representing the specific concept and its superordinate features. Thus, for example, there are units for person; for male and female; for adult and child. The busdriver is an adult male and the teacher is an adult female, etc.

In considering the task of the network, it is worth noting that there is not always a single right answer. Indeed, early on in a sentence, just after the presentation of the first constituent, there is a great deal of indeterminacy; the initial noun-phrase need not even be describing the agent of the sentence. Nevertheless, it is possible to view each constituent, as it is presented, as imposing constraints on the possible event-descriptions that might be correct. In this context, we can characterize the task of the network as being one of indicating, in response to each probe, what the range of possibilities might be, and of giving an indication by the activations that it assigns to the completions of the various probes, of its estimate of the probability associated with each.

4.3. Network architecture and processing. The architecture of the network, as shown in Figure 1, can be treated as consisting of two basic parts. One part is the actual comprehension mechanism itself, the part that reads in the constituents sequentially and updates the sentence representation; the other part is the output mechanism, that performs the probe completion task. The Sentence Gestalt units are in both parts, and form the interface between the two.

Processing occurs as follows. At the beginning of a sentence the pattern of

activation on the sentence gestalt units is set to all 0's, and the unit or units representing the first input constituent in the input pool are turned on. Activation feeds from the SG units and the input units to the hidden units in the comprehension part of the system, and from these to the SG units, where the initial SG representation of all 0's is replaced by a new pattern of activation reflecting the influence of the first constituent of the sentence. This representation is now part of the input at the next time step, when the next constituent is input in place of the first. This process continues to the end of the sentence.

Each of the units inside the network is a simple logistic processing unit; that is, the activation that a unit takes on is equal to the logistic function of its net input, where the net input is simply the sum over all connections coming to the unit of the input on each connection. The input on each connection is just the sum of the activation of the sending unit at the end of the connection, times the weight on the connection. Activations range from 0 to 1; weights are floating-point numbers initialized in a range between +/- .3, and adjusted according to the learning procedure described below.

Processing in the output network is also quite simple, and can occur at any point during or after the presentation of a sentence. The two inputs to the output network are the pattern on the SG units and the pattern on the probe units. This pattern consists of a single unit on, representing either a queried role or the queried filler. Activation feeds forward from the SG units and the probe input units to a set of hidden units and then from these to the probe output units, where the pattern is taken to represent the network's response to the probe.

4.4. Learning. Learning in the network occurs via the back-propagation learning procedure. When a probe is presented, the response to the probe can be compared to the response that would be correct for the sentence-event pair currently being processed, and the cross-entropy can be computed. Back-

propagation is used to adjust the connection strengths so as to minimize this measure (See St. John and McClelland, in press, for details).

It is important to note that the minima in this measure occur at those points where the activations of units in particular situations represent the probabilities that the units should be on in these situations. We think of the activations of the output units as representing the probability that the unit should be on. The training procedure can be seen as trying to find an ensemble of connection weight values that allow the network to get these probabilities correct.

In training the network, we followed the procedure of presenting a complete set of probes after the presentation of each constituent of each of a large number of training sentences. The complete set of probes consisted of a role probe and a filler probe for each role-filler pair in the event description for the sentence-event pair currently being processed.

This training procedure was intended to approximate the situation in which a language learner has just witnessed an event, so that he already has a description of it; and hears a sentence spoken about that event. We imagine that as the learner processes the sentence, he is continually (implicitly) asking himself, "how well does the machinery that I have for language comprehension allow me to correctly describe the event I have just witnessed." The question is posed in the form of the set of probes, and the answer is the set of response to the probes. The mismatch between the probes and the correct responses dictated by the description then serves as the basis for learning.

This procedure has two interesting characteristics. First, it does not provide the learner with any specific alignment between the constituents of the sentence and the corresponding constituents of the event description. Thus it forces the network to discover the solution to the bootstrapping problem mentioned earlier for itself. Second, the procedure requires the network to do its best at each timestep to predict all of the constituents of the event from

what is has seen so far. If learning reaches the global minimum in the error measure described above, then the activations will always reflect the best achievable estimates of the probabilities that the units should be on at each point in the processing of every sentence.

Several different runs of the model have been undertaken. The one from which we report results here involved 630,000 training trials, each involving the presentation of an independently generated sentence-event pair. Learning takes so long in part because the network is exposed to some of the lower-frequency events and contingencies only rarely. A discussion of the timecourse of acquisition is provided in St. John and McClelland (in press).


4.5. Results. After training, the model was first tested on a set of 55 randomly generated sentences that are unambiguous given the hard constraints built into the corpus. That is, although each of these sentences actually contained at least one ambiguous word or unspecified filler, the hard constraints built into the corpus were enough to allow it to respond correctly to all probes. For example, "The teacher at the soup with the utensil" is unambiguous since the only utensil that could be used for eating soup is a spoon. After presentation of each sentence, we tested the full set of probes for the role-filler pairs in the event described by the sentence. The network activated all of the correct output units more strongly than any output units it should not have activated on more than 99% of the probes.

The network was also tested specifically on several sets of sentences designed to assess its ability to handle different aspects of the comprehension task. The tasks are broken into two broad categories, having to do with role assignment on the one hand and specification of the identity of role fillers on the other. With regard to role assignment, St. John and McClelland probed with fillers from the events described by test sentences and examined the roles assigned to these fillers. The use of both syntactic and semantic constraints was examined. Thus for a sentence like "The schoolgirl stirred the kool-aid

with the spoon", semantic constraints must be used to determine that the spoon is an instrument, not an accompanier of the schoolgirl (cf "The schoolgirl stirred the kool-aid with the teacher"). In other sentences, syntactic constraints were examined. Thus for the sentence "The busdriver was given the rose by the teacher", the order of the constituents to get them with the presence of the passive marker and the preposition "by" are necessary to determine the correct role assignments of "busdriver" and "teacher", since either could play the role of agent or recipient. In tests involving 5 sentences of each of 4 types (active, passive crossed with a need to rely on semantic or syntactic constraints) all of the fillers were assigned to the correct roles. Figure 3 illustrates a passive syntactic role assignment case. Examples illustrating the other kinds of cases may be found in St. John and McClelland (in press).

For the specification of fillers, there were three distinct variants considered: The first is the straightforward resolution of word ambiguity, in which the network is asked simply to choose between two alternative quite distinct interpretations of the fillers of one or more roles. For example, in "The pitcher hit the bat with the bat", the subject, object and prepositional phrase object are all ambiguous words in the corpus, but each is sufficiently constrained by the context to yield a unique interpretation. The second variant is concept instantiation. The sentence "The teacher kissed someone" illustrates a particularly interesting case, since the someone cannot be resolved uniquely given the context but can be resolved partially. In the experience of the network, the teacher is a female, and the event generator is constrained so that kissing is always a heterosexual activity; but it occurs indiscriminately as to age. So the teacher is just as likely to kiss the pitcher (a child) or the busdriver (an adult). Thus we would expect the model to be able to identify the someone as a male but not to determine his age or whether specifically it was the pitcher or the busdriver. The third type involves what might be called "inference of implicit arguments", since in this case the sentences contained no overt indication even that there was a filler of a particular role. For example, in "The teacher at the soup", there is no instrument mentioned; but during training,

The busdriver was given the rose by the teacher.



The teacher kissed someone.




Figure 3: Activations of relevant output units in response to the indicated probes after presentation of the sentences shown.

the eating of soup always occurred with the use of a spoon, and so the spoon is inferable in this context. The model was tested with 5 different example sentences for each of the three variants. It was probed with the roles, and the output was examined to see if the correct concept units were activated. In all cases it performed correctly. In Figure 3, the output for "the teacher kissed someone" is shown, where the context partially specifies the filler (See St. John and McClelland for other examples). Here we can see that the sex, but not the age, is clearly specified. (There appears to be a slight preference for the pitcher over the busdriver. Often these slight preferences reflect the effects of specific training trials that occurred just prior to testing.)

5. How Does the Model Work?

In this section, we begin by following the time course of processing one fairly complex sentence, to give the reader a feeling for the step-by-step processing activity that occurs in the model. We then return to the questions raised at the beginning of this paper to see how the model gives very different answers to each of these questions.

The sentence we shall study is "The adult ate the steak with daintiness". The sentence is interesting, in that there are three different sources of information as to the identity of the subject. One of these is the word adult itself; the second is the fact that the adult is eating steak, since predominantly it is the male adult (the busdriver) who eats the steak; and the third is the adverb (with daintiness); in the model's experience it is only the teacher (a female) who ever eats with daintiness. As we shall see, the example illustrates the model's ability to make use of a variety of cues of varying strength, spread throughout the sentence, to identify a particular constituent.

After the presentation of each constituent (adult, ate, steak, with daintiness) we can examine the response of the network to probes assessing the fillers of the agent, action, instrument, and patient roles (See Figure 4). Later we will return to consider the pattern of activation over the SG units, which provides the representation of the whole sentence.

The adult ate the steak with daintiness.

Sentence Gestalt Activations					Role/Filler Activations				
unit	#1	#2	#3	#4		#1	#2	#3	#4
1	■	□	□	□	agent				
2	□	■	■	■	person	■	■	■	■
3	■	■	■	■	adult	■	■	■	■
4	□	□	■	■	child	□	□	□	□
5	□	□	□	■	male	□	■	■	□
6	□	□	□	□	female	□	□	□	■
7	□	□	□	□	busdriver	■	■	■	■
8	□	□	□	□	teacher	□	■	□	■
9	□	□	□	□	action				
10	□	□	□	■	ate	□	■	■	■
11	□	□	□	■	shot	□	□	□	□
12	□	□	□	□	drove(trans.)	□	□	□	□
13	□	□	□	□	drove(motiv.)	□	□	□	□
14	□	■	■	■	patient				
15	■	■	■	■	person	□	□	□	□
16	■	□	□	□	adult	□	□	□	□
17	□	■	■	■	child	□	□	□	□
18	□	□	□	■	busdriver	□	□	□	□
19	□	■	■	■	schoolgirl	□	□	□	□
20	□	□	■	■	thing	□	■	■	■
21	■	□	■	■	food	□	■	■	■
22	□	■	■	■	steak	□	■	■	■
23	□	□	□	□	soup	□	□	□	□
24	□	□	□	□	crackers	□	□	□	□
25	□	□	□	□	adverb				
26	□	□	□	□	gusto	■	■	■	■
27	■	□	■	■	pleasure	□	□	□	□
28	□	□	□	□	daintiness	□	□	□	■

Figure 4: Activation of a subset of the sentence gestalt units (on the left) and of relevant output units in response to the indicated probes (on the right) after presentation of each constituent of the sentence "The adult ate the steak with daintiness".

We consider first the response to "agent", since it is here that we see the effects of several constituents operating most clearly. After the presentation of "adult" the model takes the agent to be an adult person; there is some activation of both male and female, and of both busdriver and teacher, the only two adults in the set. There is a slight bias favoring male. Child is included to illustrate that it is not active at any point. There is little change after the presentation of the verb, since this does not really provide any constraints on the identity of the adult (the teacher and the busdriver appear equally often in sentences involving eating). The presentation of "steak", however, produces a shift in the direction of male and busdriver. This shift is reversed (though not completely) when the final constituent, "with daintiness" is presented.

For the other roles, the reader will note that the model performs in a generally sensible way. The one slight problem appears in the case of the patient. We see the activation of "steak", which was quite strong just after the presentation of the steak constituent, weaken considerably when "daintiness" is presented. We will return to a consideration of this specific aspect of the model's performance below.

Given these successes of the model, let us now ask, what kinds of answers do we get to the questions raised at the beginning of this paper when we use a model of this sort?

5.1. What is constructed when we comprehend a sentence? In this case, the answer is not "a structural description". What is constructed is a pattern of activation which permits the performance of a specific task or tasks. In this case the task is to provide a basis for completing role-filler pairs; but one can imagine a wide variety of other uses as well. Whatever the tasks were that we were called upon to use the results of comprehension to perform, a model with the general structure of the one used here could be used to learn to perform that task.

Given this, it becomes a matter of empirical research to ascertain just how a

network will choose to use its units in learning to perform the tasks that it is given to perform. We know from other connectionist research that the answers to these questions are dependent both on the specific tasks the network is asked to perform, and on the details of network architecture (Hinton, 1986; McClelland, in press). In this instance, just perusing the pattern of activation in the sentence gestalt at each successive presentation of a new input constituent, we can see two things. First, that many of the units take on graded activations, and that several of these seem only partially correlated with particular role-filler activations. This suggests that the activations of particular output units in response to particular probes are generally determined by the joint influence of a number of hidden units; thus they provide a distributed, coarse-coded representation of the role-filler information conveyed by the sentence (c.f. Hinton, McClelland, and Rumelhart, 1986).

5.2. What role do words play in the comprehension process? In the present model, as each word is presented, it changes the pattern of activation in the sentence gestalt. In this case we see each word as exerting constraints on the representation. It will be noted that these constraints can in general influence the responses to all of the probes we might present after presentation of a word. Thus the presentation of "ate" affects not only responses to probes for the action but also probes for the patient; and the presentation of steak and daintiness each influence responses to probes for the agent, the patient, and the manner. Thus a word is a clue that constrains the interpretation of the event as a whole.

The influence that a particular word will have on the comprehension process of course depends on what has already been presented. But, there is a systematic contribution that each word makes. This systematic contribution is represented by the set of connection strengths from the input unit that represents a particular word to the set of hidden units inside the comprehension part of the network.

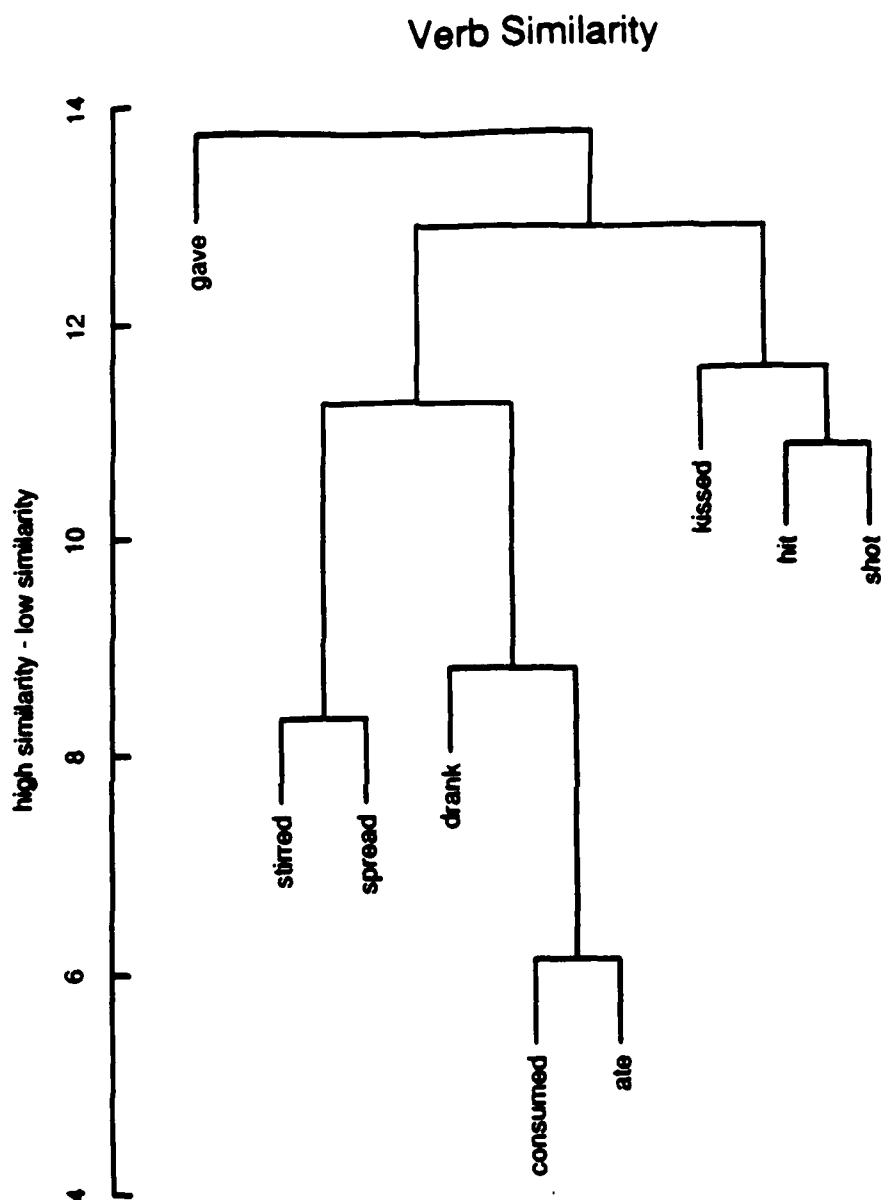


Figure 5: Cluster analysis of the weight vectors emanating from each word input unit to the hidden units in the comprehension part of the SG model, for the units representing the 11 unambiguous verbs shown. The vertical position of the horizontal bar joining two branches indicates the similarity of the leaves or branches joined.

To examine these contributions, St. John and McClelland extracted the vector of connection weights emanating from each word input unit to this first layer of hidden units. These feature vectors were then entered into a hierarchical cluster analysis; separate analyses were performed for the nouns and verbs. The analysis for the verbs (Figure 5) displays clearly that the model has captured the similarity structure among the "frames" represented by these verbs as used in our training corpus. The verb *give* is the only dative verb in the corpus, and is clustered separately from all the others. The verbs *"ate," "drank"* and *"consumed"* all take animate things as subjects and inanimate things (food) as their objects; the verbs *"stirred"* and *"spread"* each take a human subject, food as an object and a spoon or a knife as the instrument; and *hit*, *kicked*, and *kissed* are all passivisable in the corpus (unlike the food-related verbs), and all involve a patient that may be animate.

The analysis for the nouns (Figure 6) is less clear; it appears that there are two organizational principles that are both at work. Sometimes nouns cluster by meaning. Thus all the human nouns cluster separately from the rest of the nouns. However, at a finer grain, the nouns sometimes appear to cluster by co-occurrence in the same events. Thus *ice-cream* clusters with *park* because in our corpus ice-cream is eaten in the park and that is the only thing that ever happens in the park. Once again, the model appears to be picking up what might be called the frames that the nouns enter into, rather than their individual meanings per se. Of course, the details of this depend on the particular training corpus; in ordinary life, much happens in parks besides the eating of ice cream. In general it seems likely that noun-frames are much weaker than verb-frames; but to the extent that such frames do exist, they can be captured by models such as this.

5.3. How does the process of constructing a representation of a sentence occur? In the connectionist model, there is no separation of the structure sensitive rules and the lexical content of words. The process is inherently susceptible to guidance by content as well as structural information.




Figure 6: Cluster analysis of the weight vectors emanating from each word input unit to the hidden units in the comprehension part of the SG model, for the units representing the unambiguous nouns shown. The vertical position of the horizontal bar joining two branches indicates the similarity of the leaves or branches joined.

In some sense, the model represents the strongest possible alternative to a modular approach. Not only are all different sources of constraint taken into account simultaneously; the knowledge underlying each source of constraint is inextricably interwoven in the connections.

5.4. How does acquisition work? Acquisition works by a process of gradual connection strength adjustment. This is quite different from the formulation of a system of explicit rules. Certain problems are avoided right from the start, such as the question of when to form a rule, and when to simply list exceptions. However, it would certainly not be accurate to suggest that the model we have presented here is a *tabula rasa*, acquiring knowledge of language without any prior structure. Indeed, the input is parsed for the model into constituents and words; and the role-filler representation of the event descriptions and the set of concepts used in the output network are pre-determined as well. Finally, the structure of the network is pre-ordained, and tailored to the task. These features of the model were not adopted out of any belief that their adoption was necessary but simply out of a desire to establish a simple illustrative model. Just how much prior structure has to be built in, and in what way it is built in, remain basic and central issues for connectionist models in this and a number of other domains.

6. Can the PDP approach solve the problems with conventional models?

Earlier we enumerated a set of problems with conventional models. Here we consider how they are or could be solved in models of the kind we have considered here.

6.1. Conceptual guidance and rule conflicts. The problem of conceptual guidance is naturally solved by the integrated handling of both syntactic and content-based constraints on processing. The problem of rule conflicts is dealt with by the connection adjustment process. That process assigns strengths to the features so that the correct interpretations are achieved across the entire corpus.

6.2. Contextual shading as well as selection of word meaning. This characteristic of PDP models is not illustrated so clearly by the present model because of its use of local representations for concepts. We can see this kind of thing to a limited degree in such examples as "The adult ate the steak with daintiness." Though "teacher" and "female" are ultimately more active than "busdriver" and "male", the fact that it is a steak that is eaten definitely shades the activations in the network with maleness; the model seems only too natural in its ability to capture stereotypes like the one immortalized in the phrase, "real men don't eat quiche", and to use innuendo in shading its representations.

The use of local representations for concepts makes it possible to see contextual shading only in the relative degree of activation of the few superordinate feature units that were included in the model. However, this use of local representations is not inherent in the connectionist approach and we adopted this usage here only for ease of testing and to avoid building undue amounts of knowledge into the concept representations. However, an earlier model that did use distributed representations does illustrate shading effects on a grander scale (McClelland and Kawamoto, 1986). In that model, concepts were represented by fully distributed patterns. The model was trained to interpret a variety of sentences involving breaking one object with another, and all but one of the objects that could occur as the instrument shared a feature indicating that the object was hard. The one exception, the ball, was encoded as soft, and the model correctly treated it as such when it occurred in most contexts. However, when it was used to break other objects, the model shaded the representation, giving it the feature hard instead of soft; this happened just because things that break other things were typically hard, and the model became sensitive to this fact. It is worth noting that the resulting pattern was not one of the existing patterns on which the model had been trained but an extension by the model of the ensemble of possible concepts.

6.3. The similar problem with roles. The shading of concept

representations that is captured in the McClelland and Kawamoto model has been applied to roles by Touretzky and Geva (1987). The idea is simply that the set of possible roles is not some fixed set of N alternatives but an extensible set with a rich similarity structure such as is naturally captured by distributed representations.

6.4. Implied constituents. The handling of implied constituents is not a problem in the model. It is quite natural for the model to learn that events involving eating steak always involve a knife as instrument. There is no special "inference step" required to fill in the knife. This is in part a direct result of the fact that there is no prior stipulation that a particular part of the representation of the sentence corresponds to the internal reflex of each particular constituent of the sentence. It's just that events described by sentences with "ate" as the verb and "steak" as the object always involve knives as instruments. The probabilistic nature of many implied constituents is not a problem either because of the inherently graded nature of the activation process, coupled with the notion that intermediate activation values directly reflect probabilities intermediate between 0 and 1.

6.5. Combinatorial explosion or premature commitment. The model avoids combinatorial explosion by keeping multiple alternatives implicit in the single pattern of activity over the sentence gestalt. It avoids the catastrophic side-effects of premature commitment because its graded activations can be adjusted as each new constraint is introduced. In a sense it does make commitments as each new constituent is encountered, but these are not all-or-none choices but simply continuous shifts in the pattern of activation. Thus commitments made can be reversed without any backtracking. It is true that some constituents cause a more marked adjustment of the SG representation than others. These marked adjustments can be related to experimental data on reading times if we make the simple assumption that larger adjustments take longer to make. This assumption holds in systems that adjust their activations continuously (McClelland, 1979) rather than in a single time step.

We view these continuous systems as more realistic than the discrete time-step system used here; as with the use of localist representation, the use of discrete time in the illustrative example model is simply a matter of greater tractability.

6.6. The difficulty with acquisition. The use of gradual connection adjustments in the model helps it overcome some of the problems conventional approaches face in learning to interpret sentences. First, the strengths of constraints imposed by various words on the interpretation process are naturally graded and are brought gradually into balance by the connection adjustment process. Second, the solution to the bootstrapping problem emerges naturally through the exposure of the model to the statistical properties of an ensemble of sentence-event pairs. It is true that the sentence "the boy kissed the girl" could map onto the event of a boy kissing a girl in two different ways; but these alternatives are further constrained by other sentences. Thus in every sentence where the subject of the verb "kiss" is girl, there is a girl in the event and she is the agent.

We do not wish to suggest at all that the problems of acquisition are fully solved by the present model; the sentences and events are highly simplified, and the prestructuring of sentences into words and constituents, together with the prestructuring of events into role-filler pairs certainly makes things easier for the model. Our only claim is that the connectionist learning procedure we have used does have some significant advantages over rule-learning approaches. As noted above, it remains for further research to establish how much support these procedures require from pre-existing structure and how much they can induce from the environment.

7. Arguments against the PDP Approach

Several different types of arguments might be given in favor of conventional approaches and against the PDP approach to natural language. Here we consider three that seem particularly central. In all three cases, we believe that the arguments are less compelling than proponents of alternatives have alleged.

7.1. Systematicity and productivity. In their critique of connectionist models, Fodor and Pylyshyn point out that an inherent feature of the conventional approach is the fact that it accounts for the systematicity and productivity of language. They argue that connectionist models do not obviously provide an account of these facts.

Let us examine these characteristics. Systematicity refers to the fact that if a speaker can understand a sentence like "John loves the girl" and (let us say) "Bill dislikes the teacher" then he can also understand other sentences, such as "John loves the teacher", "Bill dislikes John", etc. In other words, sentences are not just isolated unanalyzed wholes but are composed of parts which can be recombined to produce other sentences that the speaker will understand.

To test the capability of a model such as ours to exhibit systematicity, we generated a new corpus, containing 10 persons and 10 actions. Each of the actions could be done by any person to any person so that there were a total of 1,000 possible events. Each could be expressed in an active or passive sentence for a total of 2,000 possible sentences.

We trained the same network described above with all but a randomly chosen 250 of the possible sentences, then, after training, we tested it on the remaining 250 sentences. A stringent accuracy criterion was adopted: A sentence was scored correct only if the unit representing the correct person or action was more active than any other unit in response to probes for the actor, action, and patient. The model got 97% of these novel sentences correct.

Now obviously this is but the first step in demonstrating that connectionist networks can exhibit systematicity. The corpus is finite, and 87.5% of it was used during training. Nevertheless, there is considerable systematicity in the model's performance.

Productivity is of course intimately linked to systematicity; it refers to the fact that we can understand many sentences that we have not actually heard

before. The experiment just described obviously addresses this point; though again, in a fairly limited way.

Other research on the productivity of connectionist networks is currently underway. Servan-Schreiber, Cleeremans, and McClelland (1988) have shown that a simple network architecture first introduced by Elman (1988) can learn to accept all of the grammatical tokens of a simple finite state language. Since in the case of this finite state language the corpus is in fact infinite, we have the first clear indication that a network can learn from finite experience to process an infinite corpus.

What remains to be established is the ability of connectionist networks to cope with languages involving long distance dependencies and embedded structures. Certainly it is reasonable to ask how well the approach taken here might be expected to extend to these more complex languages. Just to indicate that a direction exists for examining these issues, we note that the present model can easily be adapted to the processing of sentences with embedded structures. To do so, we need to enrich the query language that we use in probing the network, along with the complexity of the sentences used. One simple way to enrich the query language would simply be to probe for the third member of head-role-filler triples. Since arbitrary propositional structures can be built out of such triples, this seems like a reasonable representation language. Another possibility would be to present queries in the form of actual questions. Simulations pursuing these possibilities are underway (for some relevant demonstrations, see Mäkeläinen and Dyer, 1989).

There is an aspect of the productivity of language that appears to be better explained by our connectionist approach than by conventional approaches. This is the use of context to shade meanings of concepts as they are instantiated in particular events which may be contextually appropriate. An example of the ball from McClelland and Kawamoto illustrates this. In another case, they presented their model with the sentence "The doll moved". This

sentence was novel to the model. Among the features that the model had learned were associated with "doll" were inanimacy. However, in interpreting this sentence the model "animated" the doll. This is because, in all of the sentences that the model had been trained on, The subject of sentences of the form "X moved" were always animate. It seems to us that this interpretive liberty on the part of the model is entirely correct and appropriate, and illustrates a productivity that extends far beyond the capabilities of conventional models.

7.2. Beyond Compositionality. We have discussed two out of the three characteristics Fodor and Pylyshyn claim language has that are captured by conventional approaches. The third characteristic is compositionality: The idea that a word contributes the same thing to the meaning of all of the sentences in which it occurs. In the introduction we criticized the notion of compositionality, indicating that in fact it represents an impoverished view of the comprehension process. In our illustrative model, a word does always exert the same influence on the net input to the first set of hidden units in the comprehension part of the model. But, due to the non-linearities in the hidden units at that layer in the network, these non-linearities allow the actual impact of the word to differ greatly from context to context. The compositional contribution that a word can make in the Fodor and Pylyshyn approach can be captured; in addition, a contribution that goes beyond compositionality, encompassing context sensitivity, can be captured as well.

7.3. Lexical and syntactic modularity. We turn now to a set of considerations that arise from psychological experiments, where it is claimed that at least during some initial stage of processing, both lexical access (i.e., activation of the possible meanings associated with words) and syntactic processing (i.e., assigning attachment relations among sentence constituents) are autonomous processes. These claims run directly counter to the basic tenets of the approach that we have taken. What is the evidence?

7.3.1. *Lexical access.* In well known experiments (Swinney, 1979; Tanenhaus, Leiman, and Seidenberg, 1979) subjects listen to a spoken text containing an ambiguous word (such as BUGS) and are probed for a lexical decision immediately after the offset of the word with another word related to either meaning of the ambiguity. The oft-cited result of such experiments is the finding that decisions to words related to either meaning of the ambiguity are faster than decisions to unrelated words, indicating that both meanings are initially accessed; only later is the ambiguity resolved to fit the context so that the contextually appropriate reading is the only one that remains active.

There are two points. The first is that a recent meta-analysis (St. John, 1988) of a total of 19 studies, using both lexical decision and word naming methods, reveals that in fact there is a reliable advantage for the contextually appropriate reading, even at an immediate test. The general pattern exhibited in Figure 7 from the seminal experiment of Swinney, 1979)) is exemplary of the general pattern of the results.

The second point is that this pattern is very close to what is found in a simulation of the process of settling on an interpretation of an ambiguous word in a PDP model of the disambiguation process (Kawamoto, 1985, 1988; see Figure 7). Kawamoto's model differs from the illustrative model described here in three crucial ways. First, it uses a continuous, gradual activation process, so that units gradually settle into their final state, rather than being thrust into a state in a single step. Second, it makes use of full recurrence in the connections among the units, so that units within the same part of the system feed back on each other. Third, it does not actually simulate the full process of sentence interpretation but only considers the process of settling on a interpretation of an individual word as a joint function of contextual and phonological input. We view Kawamoto's model as an attempt to characterize the fine grain temporal processes involved in lexical access that is more coarsely approximated in the SG model.

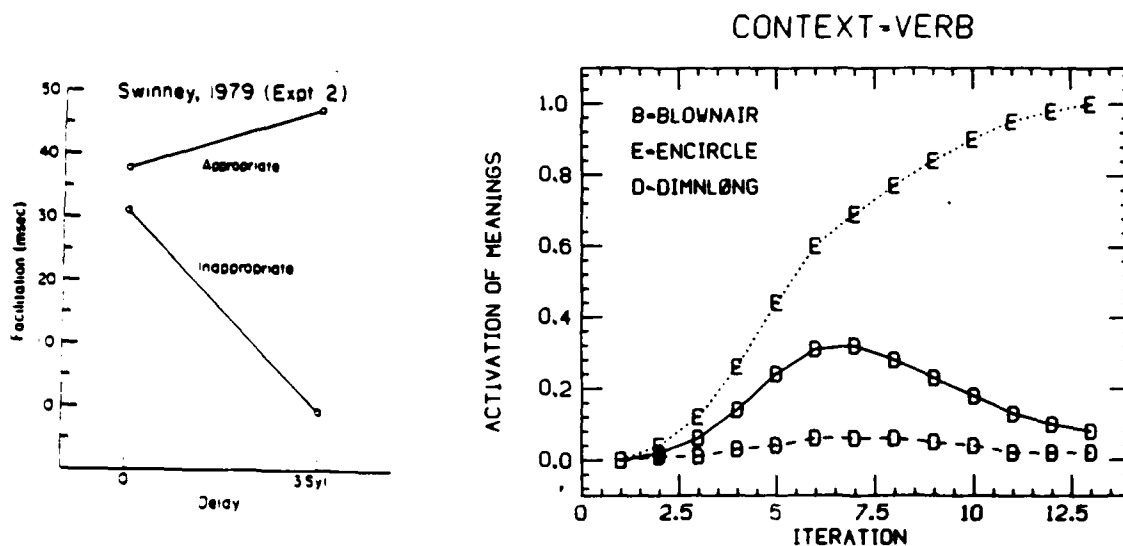


Figure 7: On the left data from Swinney, 1979; on the right, activations of meanings contextually appropriate and inappropriate meanings of ambiguous words from Kawamoto's distributed model of ambiguity resolution. Figure on the left is reprinted from McClelland, J. (Case for Interactionism); the one on the right is from Kawamoto, A. H. (1985). *Dynamic processes in the (Re)Solution of Lexical Ambiguity*. Doctoral dissertation, Department of Psychology, Brown University.

Now, Kawamoto's model most clearly does not assume that the process of accessing meaning is autonomous, in that both contextual and input-based constraints are influencing the process from the start. However, what happens in the model is that at first both of the possible meanings consistent with the input word are activated. It is only as the activation process continues that one interpretation is gradually pushed out and the other comes to dominate completely. Thus it appears that the empirical evidence is quite similar to what should be expected on a non-encapsulationist, PDP account.

7.3.2. *Autonomous syntax?* A number of studies have been reported indicating that syntactic preferences initially determine the outcome of on-line parsing processes, so that sentences in which the content eventually requires an alternative interpretation are processed more slowly than those in which the content is consistent with the syntactic bias. A variety of constructions have been examined in studies of this type. One of these is the reduced relative construction, in sentences like:

- 24. The actress sent the flowers was very pleased.
- 25. The florist sent the flowers was very pleased.

Another is the N-V-N-PP construction, as in:

- 26. The spy saw the policeman with the binoculars, but ...
- 27. The spy saw the policeman with the revolver, but ...

In the first kind of study, it is shown that subjects have difficulty processing the reduced relative clause in both cases, even though in one of the examples (the actress sent the flowers) semantic constraints are said to favor the idea that the actress would be the recipient rather than the sender of flowers as is required in the reduced relative interpretation.

Such a finding is, in our view, not particularly telling in indicating whether there is some *initial* syntactic process that favors one interpretation over the other, or whether, alternatively, there is simply a strong *weight* associated with the syntactic preference to treat a NVN sequence as actor-action-object. It certainly is the case that the initial part of the sentence:

The actress sent the flowers...

is unambiguously interpreted by native speakers as indicating that the actress is the sender not the recipient of the flowers; plausible continuations might involve a recipient (herself, perhaps?) or another clause. Thus it appears that the syntactic cues are simply overriding in this case.

In the second kind of study, the finding, as reported by Rayner, Carlson, and Frazier (1983) was that there was an advantage for sentences of the form of 26, in which the prepositional phrase is ultimately attached to the verb phrase, compared to sentences of the form of 27, in which the prepositional phrase is ultimately attached to the noun phrase. However, a series of experiments (Taraban, 1988; Taraban and McClelland, 1988; *in press*) has now established several important findings regarding this particular construction. Experiment 1 of Taraban and McClelland established three basic points. First, the materials used by Rayner et al generally had a bias such that the part of the sentence preceding the disambiguating word (revolver or binoculars, in this case) tended to favor the VP attachment of the prepositional phrase. Second, other materials are easily constructed in which this attachment preference is reversed. Third, studies of on-line processing using the word-by-word reading task developed by Just, Carpenter, and Woolley (1982) revealed that the finding reported by Rayner et al (1983) only holds with the VP attachment biased materials, and is reversed with the NP biased materials (Figure 8): With VP attachment biased materials (the Rayner, Carlson and Frazier materials), there is a reading time advantage for noun-fillers that accord with the VP attachment bias, which totals about 100 msec and is distributed over the three words following the noun-filler. However, with NP attachment biased materials (the Taraban and McClelland sentences) there is an approximately equal and opposite pattern; averaging the two types of materials, there is virtually no overall advantage for either type of attachment. Thus, the study indicates that content, rather than any general syntactic preference, appears to determine initial attachment preferences in this kind of construction.

Another experiment (Taraban and McClelland, *in press*) addressed the




Figure 8: Reading time advantage (negative numbers) or disadvantage (positive numbers) for sentences requiring a Verb-Phrase attachment of a prepositional phrase compared to matched sentences requiring a Noun-Phrase attachment. The Taraban and McClelland stimuli are biased so that subjects expect the PP to attach to the NP. The Rayner, Carlson and Frazier stimuli have the opposite bias.

question as to the source of the content-based influences on processing of the prepositional phrase. One possibility that is often considered is the idea that the verb may provide a basis for expectations about possible arguments that might influence the course of processing; these expectations could still be attributed to the workings of an autonomous syntactic process which nevertheless consulted syntactic information in the lexicon. In this experiment, Taraban and McClelland demonstrated, however, that the content of the object NP also influenced performance. For example, in sentence 28,

28. The dictator viewed the masses from the ...

29. The dictator viewed the petition from the ...

subjects expected a locative PP, attaching to the Verb, indicating the place from which the viewing was to occur; while in 29, they expected a source of the petition, attaching to the Object NP. When these expectations were violated, there was a slowdown in processing. It remains to be shown that the subject NP can also influence on-line processing. It is known from Oden (1978) that it can influence ultimate interpretations, and it seems very likely that it can influence on-line processing as well.

Another experiment 2 of Taraban and McClelland (1988) considered the possibility that the disruption in processing that is occurring in these sentences is due to specific expectations for particular fillers rather than expectations concerning the role and/or attachment of the prepositional phrases. Though a small effect for particular fillers was found, the largest effect appeared to be due to violations of expectations for the role of the prepositional phrase. Violations of expected attachment had no further disruptive effect over and above that attributable to the inevitable concomitant violation of subject's role expectations. See Taraban and McClelland (1988) for details. These findings are certainly consistent with the SG model, in that there is no separate representation of the syntactic form of a sentence; there is instead direct processing into a representation which can be used to answer questions about the roles of the participants in the event that is described by the sentence.

In summary, the evidence from the PP attachment studies seems consistent with the view that content can indeed play a role in setting up expectations for the roles played by the objects of prepositional phrases and that these expectations can govern the initial processing of these phrases as they are encountered on-line in sentence processing. Though it is very clear that syntax often exerts an over-riding influence, there is no reason to suppose on the basis of the studies reviewed here that it occupies a privileged or autonomous position in the initial processing of sentences. Instead it appears that content as well as syntax can influence initial attachment and role assignment preferences.

Further arguments against the autonomy of syntax come from the research of Crain and Steedman (1980), Altmann and Steedman (1988), and Altmann (1988). These papers argue that attachment decisions can be governed by referential processes triggered by context presented prior to the sentence containing the ambiguity. Taken together with the Taraban and McClelland results, these results help paint a general picture in which syntax is far from autonomous.

Altmann and Steedman (1988) point out that the findings on attachment ambiguity resolution are consistent with a view they call "weak interactivity", in which a syntactic module constructs all possible parses and the candidate that best satisfies all of the constraints is selected by subsequent processes sensitive to content, referential coherence, etc. They point out that such a weak interactivity account is probably not distinguishable empirically from strongly interactive accounts, in which conceptual/referential modules in the language processing system instruct modules specialized for syntactic processing.

The view taken here goes beyond even strong interactivity accounts in proposing that the syntactic and conceptual aspects of processing are in fact inextricably intertwined. Perhaps this view might best be called an integrative

as opposed to interactive account. Interactivity suggests separate systems that exert simultaneous mutual influence (c.f., Rumelhart and McClelland, 1981; McClelland, 1987), whereas in the present approach there is but a single integrated system in which syntactic and other constraints are combined in the connection weights.

7.4. Neuropsychological dissociations.

This integrative approach is actually quite different from the position one of us has taken in previous publications (McClelland, 1987). We have adopted it here, not out of any strong a priori commitment but because it has turned out to work well in capturing the phenomena considered in this paper. Indeed, the notion that there is a separate module for syntax is so ingrained in theoretical treatments of language processing that it is difficult even for us to be fully comfortable with abandoning it. But, the successes of the SG model in dealing with some of the central difficulties facing conventional approaches, coupled with the fact that the empirical evidence is beginning to favor at least some form of an interactive account, makes us feel that it is worthwhile to see if indeed there is any real basis for this implicit acceptance of some form of modularity.

In this connection, it is worth considering evidence from neuropsychology, since some of the most often-cited evidence for the view that there are separate processing systems for syntactic and conceptual information come from neuropsychological dissociations. It is generally claimed, for example, that Wernicke's aphasics have a general deficit in comprehension of word and sentence meaning which interferes with their understanding of all sentences regardless of their syntactic complexity while Broca's aphasics have a specific deficit in the ability to make use of syntactic information for comprehension. Such dissociations invite a modularist approach, in which one part of the system is specialized for use of content information and the other for the use of syntactic cues in comprehension. Could such findings possibly be consistent

with the framework considered here, in which syntactic and content-based influences on processing are inextricably intertwined in the structure of the language processing mechanism?

In fact, the notion that the difference between the Wernicke's and the Broca's aphasic can be characterized in terms of syntax and semantics is being called into question from several different vantage points. First, Milberg, Blumstein, and Dworetzky (1988) have recently reported that both Wernicke's and Broca's aphasics have differ from normals in lexical access, though the differences are complementary. Normals show a graded decrement in priming as primes are increasingly distorted, but Broca's aphasics show priming only when the prime is undistorted, and Wernicke's aphasics show priming over a wider range of distortion than normals. This suggests that Wernicke's aphasics may be suffering from something akin to undamped activation while Broca's aphasics are suffering from overdamping. Other studies suggest that Broca's and Wernicke's aphasics both show comprehension deficits, and that the deficits differ more between aphasics who speak different languages than they differ between aphasics who speak the same language. For example, Bates et al (1987) studied Broca's, Wernicke's and normal English, German, and Italian speakers. They found that within each language, Broca's and Wernicke's both show deficits in the use of morphological cues, and that the degree of preservation of the use of these cues correlated with the extent of reliance on these cues in the speaker's language. Thus Italians, who when normal rely much more on agreement and much less on word order than English speakers, showed the least impairment in the use of subject-verb agreement to mark agency, while English aphasics showed the greatest impairment. German is intermediate between the two languages in the extent of normal reliance on word order vs agreement cues, and showed an intermediate degree of disruption of the use of agreement with damage. The findings of this study are consistent with the idea that both aphasic groups show the greatest deficits in the use of cues that are relatively weaker in their native language (Bates and Wulfeck, in press; McDonald and MacWhinney, in press).

We do not mean to suggest that there is no basis at all for the idea that there may be specific dissociations of aspects of linguistic knowledge that call into question the idea that content and syntactic constraints are as fully integrated as they are in the approach that we have taken. There are several studies supporting the idea that there are particular deficits in the use of closed-class words that are restricted to Broca's and not to Wernicke's aphasics which have yet to be reconciled with the type of account suggested by the Milberg et al findings. Our only claim here is that the neuropsychological evidence is not as clear-cut as it may often appear to be, and there is room for a consideration of the idea that there may indeed be a single processing system that is simply thrown out of regulation in slightly different ways in Broca's and Wernicke's aphasia. The model we have proposed does not of course offer any insight into this differential disruption, but the model is compatible with the idea that there is a single system which uses syntax and content together to guide the language comprehension process.

8. Future Directions.

In this paper we have described an alternative to traditional models of language processing. We have tried to indicate how this alternative may allow us to solve many of the problems facing traditional approaches, and how it may provide different ways of conceptualizing basic aspects of the problem of sentence comprehension. We have also indicated that many of the arguments against the type of approach we have taken can be countered. Of course the facts are not all in, but given what is known at this time the approach seems to us to be at least as viable as any other that we know of. The model we have offered is of course far from the final word, and many problems need to be addressed our only goal has been to suggest that there may be some basis for optimism that further development of the approach might be successful.

There are several further developments which we are currently pursuing or intend to pursue shortly. First, we need to find ways of improving the rate of learning; as things stand, learning is unduly slow, especially given the small

size of the corpora that we have used in our training experiments. Second, we need to extend the framework to different kinds of output tasks. The role-filler completion task that we have used here may have several difficulties: the role-filler pair language is insufficiently structured, and the localist representation of concepts lacks the reliance on distributed representation which is one of the strengths of the PDP framework. Third, our long-term goal is to move in the direction of capturing the influence of broader, extra-sentential context on sentence processing. Ultimately, the approach will stand or fall on its ability to capture the pervasive influences of these extra-sentential factors.

References

- Altmann, G. (1988). Ambiguity, parsing strategies, and computational models. *Language and cognitive processes*, 3, 73-97.
- Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, 30, 191-238.
- Bates, E., Friederici, A., Miceli, G., & Wulfeck, B. (1985). Sentence comprehension in aphasia: A cross-linguistic study. Manuscript, University of California, San Diego.
- Bates, E., & Wulfeck, B. (in press). Crosslinguistic studies of aphasia. In B. MacWhinney & E. Bates (Eds.), *The crosslinguistic study of sentence processing*. New York: Cambridge University Press.
- Chomsky, N. (1988). Lecture presented at the University of Pittsburgh, Fall, 1988.
- Clark, H. H., & Clark, E. V. (1977). *Psychology and language: An introduction to psycholinguistics*. New York: Harcourt Brace Jovanovich, Inc.
- Crain, S., & Steedman, M. J. (1985). On not being led up the garden path: The use of context by the psychological parser. In D. Dowty, L. Karttunen, & A. Zwicky (Eds.), *Natural language parsing: Psychological, computational, and theoretical perspectives*. Cambridge, MA: Cambridge University Press.
- Elman, J. L. (1988). *Finding structure in time* (CRL Technical Report 8801). San Diego: University of California, Center for Research in Language.
- Fillmore, C. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory*. New York: Holt, Rinehart and Winston.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.
- Frazier, L. (1986). Theories of sentence processing. In J. Garfield (Ed.), *Modularity in knowledge representation and natural language processing*. Cambridge, MA: MIT Press.
- Gleitman, L. R., & Wanner, E. (1982). Language acquisition: The state of the state of the art. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: The state of the art*. Cambridge, MA: Cambridge University Press.
- Hinton, G. E. (1986). Learning distributed representations of concepts. *Proc. Eighth Annual Conference of the Cognitive Science Society*, Amherst, MA.
- Hinton, G. E. (1987). *Connectionist learning procedures* (Technical Report

CMU-CS-87-115). Pittsburgh, PA: Carnegie Mellon University, Department of Computer Science.

- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart, J. L. McClelland, & the PDP research group. (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume I.* Cambridge, MA: Bradford Books.
- Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, 111.
- Kawamoto, A. H. (1985). *Dynamic processes in the (re)resolution of lexical ambiguity.* Doctoral dissertation, Department of Psychology, Brown University.
- Kawamoto, A. H. (1988). Distributed representations of ambiguous words and their resolution in a connectionist network. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence.* San Mateo, CA: Morgan Kaufmann Publishers, Inc.
- Marcus, M. P. (1980). *A theory of syntactic recognition for natural language.* Cambridge, MA: MIT Press.
- McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, 86, 287-330.
- McClelland, J. L. (1987). The case for interactionism in language processing. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading.* London: Erlbaum. 1-36.
- McClelland, J. L. (in press). Parallel distributed processing: Implications for cognition and development. In Morris, R. (Ed.), *Parallel distributed processing: implications for psychology and neurobiology.* Oxford University Press.
- McClelland, J. L., & Kawamoto, A. H. (1986). Mechanisms of sentence processing: Assigning roles to constituents. In J. L. McClelland, D. E. Rumelhart, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume II.* Cambridge, MA: Bradford Books.
- McDonald, J., & MacWhinney, B. (in press). Maximum likelihood models for sentence processing. In B. MacWhinney & E. Bates (Eds.), *The crosslinguistic study of sentence processing.* New York: Cambridge University Press.
- Miikkulainen, R., & Dyer, M. G. (1989). *A modular neural network architecture*

for sequential paraphrasing of script-based stories (Technical Report UCLA-AI-89-02). Los Angeles: University of California, Artificial Intelligence Laboratory, Computer Science Department.

- Milberg, W., Blumstein, S., & Dworetzky, B. (1988). Phonological processing and lexical access in aphasia. *Brain and Language*, 34, 279-293.
- Oden, G. (1978). Semantic constraints and judged preference for interpretations of ambiguous sentences. *Memory and Cognition*, 6, 26-37.
- Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behavior*, 22, 358-374.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Parallel distributed processing models of schemata and sequential thought processes. In J. L. McClelland, D. E. Rumelhart, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume II*. Cambridge, MA: Bradford Books.
- Servan-Schreiber, D., Cleeremans, A., & McClelland, J. L. (1988). *Encoding sequential structure in simple recurrent networks* (CMU-CS-88-183). Pittsburgh, PA: Carnegie Mellon University, Computer Science Department.
- Schank, R. C. (1981). Language and memory. In D. A. Norman (Ed.), *Perspectives on cognitive science*. Norwood, NJ: Ablex Publishing Corporation.
- St. John, M. F. (1988). Hitting the right pitch: A meta-analysis of the processing of ambiguous words in context. Manuscript.
- St. John, M. F., & McClelland, J. L. (in press). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*.
- Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18, 645-659.
- Tanenhaus, M. K., Leiman, J. M., & Seidenberg, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*, 18, 427-440.
- Taraban, R. (1988). *Content-based expectations: One source of guidance for syntactic attachment and thematic role assignment in sentence processing*. Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA.
- Taraban, R., & McClelland, J. L. (1988). Constituent attachment and thematic

role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language*, 27, 597-632.

Taraban, R., & McClelland, J. L. (in press). Parsing and Comprehension. A multiple-constraint view. In K. Rayner, M. Balota, & I. Flores D'Arcais (Eds.), *Processes in reading comprehension*.

Touretzky, D. S., & Geva, S. (1987). A distributed connectionist representation for concept structures. Paper presented to the 9th Annual Conference of the Cognitive Science Society, Seattle, WA.

van Gelder, T. (in press). Compositionality: Variation on a classical theme. *Cognitive Science*.